

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Appellants:	Samuel FINEBERG et al.	§	Confirmation No.:	5258
		§		
Serial No.:	10/808,138	§	Group Art Unit:	2185
		§		
Filed:	March 24, 2004	§	Examiner:	Daniel Tsui
		§		
For:	Communication-Link-	§	Docket No.:	200314538-1
	Attached Persistent	§		
	Memory System	§		

**APPEAL BRIEF**

**Mail Stop Appeal Brief – Patents**

Commissioner for Patents  
PO Box 1450  
Alexandria, VA 22313-1450

Date: April 17, 2008

Sir:

Appellants hereby submit this Appeal Brief in connection with the above-identified application. A Notice of Appeal is being electronically filed concurrently herewith.

**TABLE OF CONTENTS**

I.	REAL PARTY IN INTEREST .....	3
II.	RELATED APPEALS AND INTERFERENCES.....	4
III.	STATUS OF THE CLAIMS .....	5
IV.	STATUS OF THE AMENDMENTS.....	6
V.	SUMMARY OF THE CLAIMED SUBJECT MATTER.....	7
VI.	GROUND OF REJECTION TO BE REVIEWED ON APPEAL.....	9
VII.	ARGUMENT.....	10
A.	Section 103 Rejections over Zhang and RDMA.....	10
1.	Claims 1-2 and 8.....	10
2.	Claims 10 and 11 .....	13
B.	Section 103 Rejections over Zhang, RDMA and Golding .....	16
1.	Claims 3 and 4 .....	16
C.	Section 103 Rejections over Zhang, RDMA, Golding and Olson .....	16
1.	Claim 5 .....	16
D.	Section 103 Rejections over Zhang, RDMA and Garg .....	16
1.	Claims 6 and 7 .....	16
E.	Section 103 Rejections over Zhang, RDMA and AIPI .....	16
1.	Claim 9 .....	16
F.	Conclusion.....	16
VIII.	CLAIMS APPENDIX.....	18
IX.	EVIDENCE APPENDIX .....	21
X.	RELATED PROCEEDINGS APPENDIX .....	22

**I. REAL PARTY IN INTEREST**

The real party in interest is Hewlett-Packard Development Company, L.P. (HPDC), a Texas Limited Partnership, having its principal place of business in Houston, Texas. The Assignment from the inventors to HPDC was recorded on March 24, 2004, at Reel/Frame 015145/0850.

**Appl. No. 10/808,138**  
**Appeal Brief dated April 17, 2008**  
**Reply to final Office action of February 26, 2008**

**II. RELATED APPEALS AND INTERFERENCES**

Appellants are unaware of any related appeals or interferences.

**III. STATUS OF THE CLAIMS**

Originally filed claims: 1-46.  
Claim cancellations: 12-46.  
Added claims: None.  
Presently pending claims: 1-11.  
Presently appealed claims: 1-11.

**Appl. No. 10/808,138**  
**Appeal Brief dated April 17, 2008**  
**Reply to final Office action of February 26, 2008**

**IV. STATUS OF THE AMENDMENTS**

No claims were amended after the final Office action dated February 26, 2008.

## **V. SUMMARY OF THE CLAIMED SUBJECT MATTER**

The specification is directed to a communication-link-attached persistent memory unit.<sup>1</sup> At least some of the illustrative embodiments a systems as in claim 1:

1. A persistent memory access system, comprising:
  - a primary network persistent memory unit (nPMU)<sup>2</sup> comprising:
    - a network interface communicatively coupled to at least one client processor node over a remote direct memory access (RDMA) enabled communication system;<sup>3</sup> and
    - a primary region coupled to the network interface and configured to store information, the primary region is assigned to a client process running on the client processor node and stores information received from the client process;<sup>4</sup>
  - said primary nPMU executes single byte RDMA requests directed to the primary region, the single byte RDMA requests received through the network interface;<sup>5</sup>
  - a mirror nPMU<sup>6</sup> comprising:
    - a network interface communicatively coupled to the at least one client processor node and the primary nPMU over the remote direct memory access (RDMA) enabled communication system;<sup>7</sup> and

---

<sup>1</sup> Specification Title.

<sup>2</sup> The application was filed by a different firm than is now currently prosecuting the application. In order to ensure consistent citations to the Specification, in this Appeal Brief citations to the Specification are made with regard to the application as published at U.S. Pub. No. 2005/0216552A2. The primary nPMU feature finds support in the published application on page 6, paragraph [0055], lines 1-4. A shorthand notation for citations takes the form ([page], [paragraph], lines [lines within the paragraph]). The citation of this footnote in the shorthand form reads (6, [0055], lines 1-4). *See also*, Figure 8, element 802.

<sup>3</sup> (6, [0055], lines 4-6), Figure 8, elements 806, 808; (6, [0056], lines 1-6), Figure 8, elements 802, 804.

<sup>4</sup> (7, [0062], lines 1-17), Figure 9, element 906.

<sup>5</sup> (3, [0034], lines 15-18).

<sup>6</sup> (6, [0055], lines 1-4), Figure 8, element 804.

<sup>7</sup> (6, [0055], lines 4-6), Figure 8, elements 806, 808; (6, [0056], lines 1-6), Figure 8, elements 802, 804.

a mirror region coupled to the network interface of the mirror nPUM and configured to store information, the mirror region is assigned to the client process and stores the information received from the client process.<sup>8</sup>

Other illustrative embodiments are systems as in claim 10, having all the limitations of claim 1, and further reciting:

10. The system of claim 1, further comprising a persistent memory manager (PMM) communicatively coupled to the client processor node via the communication system, wherein the PMM responds to a request from the client process for an assignment of persistent memory, and wherein the PMM causes the primary nPMU to create the primary region and causes the mirror nPMU to create the mirror region.<sup>9</sup>

---

<sup>8</sup> (7, [0062], lines 1-17), Figure 9, element 906.

<sup>9</sup> (6, [0057], lines 1-10), Figure 8, element 822.



**VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL**

Whether claims 1-2, 8 and 10-11 are obvious under 35 USC § 103(a) over Zhang et al. (U.S. Pat. No. 7,251,713, Zhang) in view of the publication titled "An Overview of RDMA over IP" (hereafter RDMA).

Whether claims 3 and 4 are obvious under 35 USC § 103(a) over Zhang, RDMA and Golding (U.S. Pat. No. 6,477,617).

Whether claim 5 is obvious under 35 USC § 103(a) over Zhang, RDMA, Golding and Olson (U.S. Pat. No. 5,479,628).

Whether claims 6 and 7 are obvious under 35 USC § 103(a) over Zhang, RDMA and Garg (U.S. Pat No. 7,266,645)

Whether claim 9 is obvious under 35 USC § 103(a) over Zhang, RDMA and AIPI (IEEE Dictionary).

## VII. ARGUMENT

### A. Section 103 Rejections over Zhang and RDMA

#### 1. Claims 1-2 and 8

Claims 1-2 and 8 stand rejected as allegedly anticipated by Frank. Claim 1 is representative of this grouping of claims. The grouping should not be construed to mean the patentability of any of the claims may be determined in later actions (*e.g.*, actions before a court) based on the groupings. Rather, the presumption of 35 USC § 282 shall apply to each of these claims individually.

The Office action of February 26, 2008 fails to make a *prima facie* case of obviousness regarding representative claim 1. In particular, claim 1 recites:


- a primary network persistent memory unit (nPMU) comprising:
  - a network interface communicatively coupled to at least one client processor node over a remote direct memory access (RDMA) enabled communication system; and
  - ...
- a mirror nPMU comprising:
  - a network interface communicatively coupled to the at least one client processor node and the primary nPMU over the remote direct memory access (RDMA) enabled communication system... .

Even if hypothetically assumed that Zhang and RDMA teach the primary region, memory region and RDMA aspects (which Applicants do not admit), the Office action fails to cite to any portion Zhang or RDMA for the claimed network interfaces. For this reason alone the rejections should be overturned and the claims set for issue.

Zhang is directed to a system and method to transport data snapshots.<sup>10</sup> Zhang's Figure 3 is reproduced immediately below for convenience of this discussion.

---

<sup>10</sup> Zhang Title.



In particular, Zhang describes a system having a primary server 402 and a secondary server 406.<sup>11</sup> The primary server 402 has a storage device 404 controlled by source controller 414, and the secondary server 406 has a storage device 408 controlled by target controller 416.<sup>12</sup> The servers 402 and 406 correspond to the computer systems 102-105 of Zhang's Figure 1, the controller 414 corresponds to the primary controllers 108-109 of Zhang's Figure 1, and the target controller 416 corresponds to the secondary controllers 112-115 of Zhang's Figure 1.

Still referring to Zhang's Figure 3 above, a backup application 301 executing in server 402 issues a backup request to its controller 414, and in response the controller 414 makes a mirror copy data in storage device 404, with the mirror copy likewise residing (initially) in storage device 404.<sup>13</sup> Another program executing on the server 402, the VSS Application 305, determines a secondary server (in this example server 406 and particularly storage device 408)

<sup>11</sup> Zhang Col. 6, lines 28-32.

<sup>12</sup> *Id.*

<sup>13</sup> Zhang Col. 6, lines 49-60.

to which to send the mirror copy.<sup>14</sup> Once the mirror copy is transferred or otherwise accessible to the storage device 408 of server 406, the data can then be backed-up without adversely impacting operation of server 402.<sup>15</sup>

A computer or server's performance, however, will be affected by the mirroring and backup process during the backup window because numerous commands need to communicate back and forth with a storage controller coupled to the storage device. ... Thus, system administrators may prefer to utilize additional servers to perform the backup process while the primary server continues to perform its intended function.<sup>16</sup>

The storage devices with which Zhang is concerned are large block-level (*i.e.*, multiple-bytes in every read/write) data storage devices.

Each of the data storage units may include a plurality of data storage devices, such as a redundant arrays of inexpensive disks (RAID). Other types of storage media and devices may also be used, such as optical disks, optical tape, magnetic tape, magnetic diskettes, solid state memory, or any other suitable storage medium.<sup>17</sup>

Thus, not only is Zhang silent as to direct memory access (DMA) or remote (RDMA), such byte-level granularity runs contrary to Zhang's intended purpose of large file-level mirroring and subsequent back-up.<sup>18</sup>

Applicants respectfully submit that Zhang and RDMA are not properly considered together. The Manual of Patent Examining Procedures (MPEP) provides guidance with respect to when references can be considered together.

**THE PROPOSED MODIFICATION CANNOT RENDER THE PRIOR  
ART UNSATISFACTORY FOR ITS INTENDED PURPOSE**

If the proposed modification would render the prior art unsatisfactory for its intended purpose, then there is no suggestion or motivation to make the proposed modification.<sup>19</sup>

---

<sup>14</sup> Zhang Col. 6, line 6 through Col. 7, line 8.

<sup>15</sup> Zhang Col. 7, lines 9-16.

<sup>16</sup> Zhang Col. 1, lines 48-58.

<sup>17</sup> Zhang Col. 3, lines 47-52.

<sup>18</sup> Zhang Col. 5, lines 34-38.

<sup>19</sup> MPEP 8<sup>th</sup> Ed., Rev. 6, August 2005, §2143.01(V), p. 2100-137.

Zhang is directed to large, file-level mirroring and subsequent backup. If Zhang's file-level mirroring is modified to operate at byte-level granularity, the ability of Zhang's computer or servers to mirror, transfer the mirror copy and backup the mirror copy without adverse effect is severely diminished if not totally destroyed by the additional overhead of transferring file-level data across a communication path geared for byte-level granularity. For this additional reason the rejection should be overturned and the claim set for issue.

Relatedly, the MPEP further admonishes:

THE PROPOSED MODIFICATION CANNOT CHANGE THE  
PRINCIPLE OF OPERATION OF A REFERENCE

If the proposed modification or combination of the prior art would change the principle of operation of the prior art invention being modified, then the teachings of the references are not sufficient to render the claims *prima facie* obvious.<sup>20</sup>

Here again, Zhang is directed to large, file-level mirroring and subsequent backup. If Zhang's large, file-level mirroring is modified to operate at byte-level granularity, then Zhang is not able to implement the RAID-type functionality.

Based on the foregoing, Applicants respectfully request that the rejections of this grouping be reversed, and the claims set for issue.

## **2. Claims 10 and 11**

Claims 10 and 11 stand rejected as allegedly anticipated by Frank. Claim 10 is representative of this grouping of claims. The grouping should not be construed to mean the patentability of any of the claims may be determined in later actions (*e.g.*, actions before a court) based on the groupings. Rather, the presumption of 35 USC § 282 shall apply to each of these claims individually.

---

<sup>20</sup> MPEP 8<sup>th</sup> Ed., Rev. 6, August 2005, §2143.01(VI), p. 2100-138.

The Office action of February 26, 2008 fails to make a *prima facie* case of obviousness regarding representative claim 10. In particular, claim 10 requires, by virtue of claim 1 from which it depends:

- a primary network persistent memory unit (nPMU) comprising:
  - a network interface communicatively coupled to at least one client processor node over a remote direct memory access (RDMA) enabled communication system; and
  - ...
- a mirror nPMU comprising:
  - a network interface communicatively coupled to the at least one client processor node and the primary nPMU over the remote direct memory access (RDMA) enabled communication system...

Even if hypothetically assumed that Zhang and RDMA teach the primary region, memory region and RDMA aspects (which Applicants do not admit), the Office action fails to cite to any portion Zhang or RDMA for the claimed network interfaces. For this reason alone the rejections should be overturned and the claims set for issue.

Further, Applicants respectfully submit that Zhang and RDMA are not properly considered together.<sup>21</sup> Zhang is directed to large, file-level mirroring and subsequent backup. If Zhang's file-level mirroring is modified to operate at byte-level granularity, the ability of Zhang's computer or servers to mirror, transfer the mirror copy and backup the mirror copy without adverse effect is severely diminished if not totally destroyed by the additional overhead of transferring file-level data across a communication path geared for byte-level granularity. For this additional reason the rejection should be overturned and the claim set for issue.

---

<sup>21</sup> MPEP 8<sup>th</sup> Ed., Rev. 6, August 2005, § 2143.01(V), p. 2100-137.

Relatedly, considering Zhang with RDMA changes the principle of operation of Zhang.<sup>22</sup> Zhang is directed to large, file-level mirroring and subsequent backup. If Zhang's file-level mirroring is modified to operate at byte-level granularity, then Zhang is not able to implement the RAID-type functionality.

Further still, representative claim 10 recites:

10. The system of claim 1, further comprising a persistent memory manager (PMM) communicatively coupled to the client processor node via the communication system, wherein the PMM responds to a request from the client process for an assignment of persistent memory, and wherein the PMM causes the primary nPMU to create the primary region and causes the mirror nPMU to create the mirror region.

The Office action attempts to rely on Zhang's storage controllers 108 and 112 for the claimed persistent memory manager; however, the Office action already relies on storage controllers 108 and 112 as the claimed primary and mirror persistent memory units. Thus, if the storage controllers 108 and 112 are the claimed persistent memory manger, then the Zhang and RDMA fail to teach or suggest the claimed persistent memory units. If the storage controllers 108 and 112 are the claimed persistent memory units, then Zhang and RDMA fail to teach or suggest the persistent memory manager. Stated otherwise, Zhang and RDMA clearly fail to teach a device that "causes the primary nPMU to create the primary region and causes the mirror nPMU to create the mirror region." For this additional reason the rejections should be reversed and the claims set for issue.

Zhang is silent as to an ability or inability of storage controllers 108 and 112 to allocate space. Any inherent teaching of Zhang with regard to allocations could only be that each storage controller allocates space only on its directly attached storage devices.<sup>23</sup> Thus, even if hypothetically proper to rely on one of Zhang's storage controller 108 or 112 for the claimed persistent memory

---

<sup>22</sup> MPEP 8<sup>th</sup> Ed., Rev. 6, August 2005, § 2143.01(VI), p. 2100-138.

<sup>23</sup> See, e.g., Zhang's Figure 1.

manager (which Applicants do not admit), Zhang and RDMA still fail to teach a persistent memory manager that “causes the primary nPMU to create the primary region **and** causes the mirror nPMU to create the mirror region” as neither controller would have the ability to allocate on non-locally coupled storage devices.

Based on the foregoing, Applicants respectfully request that the rejections of this grouping be reversed, and the claims set for issue.

**B. Section 103 Rejections over Zhang, RDMA and Golding**

**1. Claims 3 and 4**

Claims 3 and 4 are allowable for the same reasons as discussed in Section VII(A)(1).

**C. Section 103 Rejections over Zhang, RDMA, Golding and Olson**

**1. Claim 5**

Claim 5 is allowable for the same reasons as discussed in Section VII(A)(1).

**D. Section 103 Rejections over Zhang, RDMA and Garg**

**1. Claims 6 and 7**

Claims 6 and 7 are allowable for the same reasons as discussed in Section VII(A)(1).

**E. Section 103 Rejections over Zhang, RDMA and AIPI**

**1. Claim 9**

Claim 9 is allowable for the same reasons as discussed in Section VII(A)(1).

**F. Conclusion**

For the reasons stated above, Appellants respectfully submit that the Examiner erred in rejecting all pending claims. It is believed that no extensions of time or fees are required, beyond those that may otherwise be provided for in documents accompanying this paper. However, in the event that additional extensions of time are necessary to allow consideration of this paper, such



**Appl. No. 10/808,138**  
**Appeal Brief dated April 17, 2008**  
**Reply to final Office action of February 26, 2008**

extensions are hereby petitioned under 37 C.F.R. § 1.136(a), and any fees required (including fees for net addition of claims) are hereby authorized to be charged to Hewlett-Packard Development Company's Deposit Account No. 08-2025.

Respectfully submitted,

/mes/

Mark E. Scott  
PTO Reg. No. 43,100  
CONLEY ROSE, P.C.  
(512) 610-3410 (Phone)  
(512) 610-3456 (Fax)  
ATTORNEY FOR APPELLANTS

HEWLETT-PACKARD COMPANY  
Intellectual Property Administration  
Legal Dept., M/S 35  
P.O. Box 272400  
Fort Collins, CO 80527-2400

**VIII. CLAIMS APPENDIX**

1. (Previously presented) A persistent memory access system, comprising:  
a primary network persistent memory unit (nPMU) comprising:
  - a network interface communicatively coupled to at least one client processor node over a remote direct memory access (RDMA) enabled communication system; and
  - a primary region coupled to the network interface and configured to store information, the primary region is assigned to a client process running on the client processor node and stores information received from the client process;said primary nPMU executes single byte RDMA requests directed to the primary region, the single byte RDMA requests received through the network interface;  
a mirror nPMU comprising:
  - a network interface communicatively coupled to the at least one client processor node and the primary nPMU over the remote direct memory access (RDMA) enabled communication system; and
  - a mirror region coupled to the network interface of the mirror nPUM and configured to store information, the mirror region is assigned to the client process and stores the information received from the client process.
2. (Previously presented) The system of claim 1 wherein the primary nPMU and the mirror nPMU are physically separate units and are characterized by separate fault domains.
3. (Original) The system of claim 1, wherein the primary region comprises a plurality of primary virtual addresses corresponding a plurality of physical locations where the information is stored in the primary region, and wherein the mirror region comprises a plurality of mirror virtual addresses corresponding to

another plurality of physical locations where the information is stored in the mirror region.

4. (Original) The system of claim 3, wherein the primary nPMU is configured to translate between one of the primary virtual addresses and a corresponding client address associated with the information, and wherein the mirror nPMU is configured to translate between one of the mirror virtual addresses and the corresponding client address associated with the information.

5. (Original) The system of claim 4, further comprising a base pointer, the base pointer corresponding to a difference in the primary virtual address and the corresponding client address such that the primary nPMU translates, and wherein the base pointer further corresponds to a difference in the mirror virtual address and the corresponding client address such that the mirror nPMU translates.

6. (Original) The system of claim 1, further comprising metadata residing in the primary nPMU and the mirror nPMU, wherein the metadata identifies the primary region and the mirror region assigned to the client process.

7. (Original) The system of claim 6, further comprising an attribute cache, the attribute cache residing in the client processor node and corresponding to at least some attribute information in the metadata.

8. (Original) The system of claim 1, further comprising a persistent memory unit (PMU) library residing in the client processor node, wherein the PMU library comprises at least a first function configured to permit an executing process to directly write the information to the primary region and the mirror region, and a second function configured to permit the executing process to directly read the information from a selected one of the primary region or the mirror region.

9. (Original) The system of claim 8, further comprising an application program interface (API) residing in the client processor node, wherein the API causes the client process to access the functions of the PMU library in response to a request by the client process to access the information.

10. (Original) The system of claim 1, further comprising a persistent memory manager (PMM) communicatively coupled to the client processor node via the communication system, wherein the PMM responds to a request from the client process for an assignment of persistent memory, and wherein the PMM causes the primary nPMU to create the primary region and causes the mirror nPMU to create the mirror region.

11. (Original) The system of claim 1, further comprising a persistent memory manager (PMM) communicatively coupled to the client processor node via the communication system, wherein the PMM responds to a request from the client process to end an assignment of persistent memory, and wherein the PMM causes the primary nPMU to delete the primary region and causes the mirror nPMU to delete the mirror region.

12.-46. (Cancelled).

**Appl. No. 10/808,138**  
**Appeal Brief dated April 17, 2008**  
**Reply to final Office action of February 26, 2008**

**IX. EVIDENCE APPENDIX**

None.

**Appl. No. 10/808,138**  
**Appeal Brief dated April 17, 2008**  
**Reply to final Office action of February 26, 2008**

**X. RELATED PROCEEDINGS APPENDIX**

None.